




WHITEPAPER

CYBER INSURANCE IN THE AGE OF AI PHISHING

Why Underwriters Should Care About Phishing-Resistant Authentication

Davis Hake, Luke O'Grady

OCTOBER 2026


Contents

I.	Executive Summary.....	3
II.	Introduction.....	7
III.	AI-Enabled Phishing and the Changing Threat Landscape.....	8
IV.	How Does Phishing-Resistant MFA Reduce an Organization's Cybersecurity Risk and Financial Exposure? ..	12
V.	A Cyber Insurance Underwriter's Perspective on Responding to AI-Enabled Phishing.....	15
	Are Underwriters Aware of the Risks Posed by AI-Enabled Phishing?.....	15
	Are Underwriters Aware of the Differences in Effectiveness Between Legacy and Phishing-Resistant MFA? ..	16
	How Are Underwriters Assessing Phishing-Related Cybersecurity Risk Before Binding Coverage?.....	17
VI.	Current State of Adoption of Phishing-Resistant MFA.....	18
VII.	Best Practices for Deploying Phishing-Resistant MFA.....	21
VIII.	Recommendations for Underwriters.....	23
	Improve Questions for Cyber Insurance Applications.....	23
	Help Insureds Improve During the Policy Period.....	26
	Capture Data on the Underlying Causes of Cyber Insurance Claims.....	27
IX.	Future-Proofing for Agentic AI.....	27
X.	Conclusion.....	30
	Technical Annex. Legacy vs. Phishing-Resistant MFA.....	32
	What is Legacy MFA?.....	32
	What is Phishing-Resistant MFA?.....	35

I. Executive Summary

Generative AI has transformed phishing, a long-standing and persistent form of social engineering, from a relatively crude and error-prone tactic into a fast, scalable, and highly convincing attack method. Attackers can now create personalized lures, produce polished and natural-sounding language, and deploy interactive chatbots that manipulate victims in real time.

Recent data from both industry and government suggests that AI-enabled attacks are becoming both more common and more costly. IBM's *Cost of a Data Breach Report 2026* found that more than one in four malicious breaches were AI-enabled, representing a 56% increase from the previous year. These breaches cost organizations an average of \$6 million — approximately \$1 million more than breaches in which AI use was not reported.¹ Similarly, the U.S. Federal Bureau of Investigation's (FBI) Internet Crime Complaint Center (IC3) stated in its 2025 *Internet Crime Report* that it had received more than 22,000 complaints referencing AI-related activity in 2025 alone, with adjusted losses exceeding \$893 million.²



Although claims data rarely identifies AI as the direct root cause of an incident, cyber insurers are already seeing its downstream effects in the growing frequency and severity of social-engineering losses. Coalition reported that Business Email Compromise (BEC) and Funds-Transfer Fraud (FTF) accounted for 58% of all claims in 2025. Of those FTF claims, 71% were directly caused by social engineering, while 52% originated as a BEC, with an average loss of \$112,000.³ Travelers similarly

¹<https://www.ibm.com/forms/mkt-whitepaper-90d12>

²https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf

³<https://www.coalitioninc.com/claims-report/2026>

reported that BEC and social-engineering fraud have represented nearly half of its cyber claims over the past five years.⁴ Meanwhile, Resilience found that claims tied to phishing, social engineering, and FTF surged from 18% in 2024 to 85% in the first half of 2026. “Attackers have used AI to sharpen social engineering across the board, from emails to voice deepfakes, and defensive tactics have to keep pace,” the carrier notes in its *2026 Midyear Claims Report*.⁵

While legacy MFA such as SMS, OTP, and push notifications were considered as good enough, they are highly susceptible to phishing and even more so in the case of AI-driven phishing attacks. To better protect themselves, organizations should adopt phishing-resistant MFA built on FIDO2 standards — such as passkeys stored on security keys — or public key infrastructure (PKI).

Research supports the effectiveness of these controls. A study by Forrester and Yubico found that organizations deploying FIDO2 hardware security keys effectively eliminated phishing and credential-theft risk while achieving a 265% return on investment.⁶ Similarly, the Marsh McLennan Cyber Risk Intelligence Center's August 2025 report, titled *Cybersecurity Signals: Connecting Controls and Incident Outcomes*, found that phishing-resistant MFA was associated with a 2.25 times greater reduction in breach likelihood than that associated with legacy MFA.⁷

Recognizing these benefits, governments and regulators are increasingly directing organizations toward phishing-resistant authentication. However, these efforts have not yet produced the level of adoption needed to materially reduce cyber risk and the resulting financial exposure. This creates a significant challenge for cyber insurers, which ultimately underwrite much of that exposure.

Cyber insurers have influenced security practices before. As ransomware losses increased, they played a pivotal role in moving the market from passwords to MFA. Interviews conducted by the Center for Cybersecurity Policy and Law indicate that cyber insurers now recognize that many of the legacy MFA

⁴https://info.corvusinsurance.com/hubfs/Ransomware_Reports/Q2_25_Travelers_Cyber_Threat_Report.pdf

⁵<https://cyberresilience.com/wp-content/uploads/2026/08/2026-Midyear-Cyber-Risk-Report.pdf>

⁶<https://tei.forrester.com/go/yubico/yubikeys/?lang=en-us>

⁷<https://www.marsh.com/en/services/cyber-risk/insights/cybersecurity-signals.html>

methods they previously promoted are no longer sufficient. This paper therefore calls on insurers to help drive the next transition from legacy MFA to phishing-resistant MFA in the following ways:

First, underwriters should modernize the questions they ask on insurance applications. The objective is not to turn every application into a full identity and access management (IAM) audit. Rather, insurers should replace broad, low-signal questions with more specific questions that reveal how authentication controls operate in practice:

- ▶ Does the organization require MFA for email, remote access, cloud applications, and privileged accounts?
- ▶ Which authentication methods are permitted, including SMS, email one-time passwords, voice calls, push notifications, time-based one-time passwords, passkeys, FIDO2/WebAuthn security keys, certificate-based authentication, or other methods?
- ▶ Is phishing-resistant MFA required for administrators, executives, finance personnel, IT and help-desk users, developers with production access, and other high-risk users?
- ▶ Are payment changes and other high-risk financial requests verified through an out-of-band channel and subject to dual approval?
- ▶ Are weaker fallback methods, such as SMS, email one-time passwords, voice calls, or security questions allowed for accounts otherwise protected by stronger authentication?
- ▶ Can users bypass MFA through help-desk resets, account recovery, emergency access procedures, or legacy authentication protocols?
- ▶ Are vendors, contractors, and third-party service providers subject to the same authentication requirements as employees?
- ▶ Does the organization maintain separate administrative accounts for privileged activities and/or only provide privileged permissions when necessary?
- ▶ How are privileged accounts issued, reviewed, suspended, and revoked when employees or contractors change roles or leave the organization?
- ▶ Is phishing-resistant MFA enforced through conditional-access policies, equivalent technical controls, or identity-provider configuration?

- ▶ Does the organization monitor for anomalous logins, authenticator changes, impossible travel, token theft, Open Authorization (OAuth) abuse, suspicious account recovery, and dormant privileged accounts?
- ▶ Are identity controls reassessed after mergers, cloud migrations, identity-provider changes, major incidents, or the adoption of AI agents?

Second, underwriters should use the policy period to help insureds strengthen their identity posture.

Potential measures include education for brokers and insureds, partnerships with systems integrators and managed security providers, premium credits or improved terms for verified phishing-resistant MFA deployments, and hardware security key pilot programs for privileged or high-risk users.

Third, underwriters should develop a more consistent claims taxonomy that identifies the specific identity-control failure behind a loss. Rather than recording only the ultimate claim outcome, insurers should distinguish between failures such as phishable MFA, push fatigue, token theft, and account-recovery bypass. More precise claims data would improve underwriting decisions and help the market identify emerging control weaknesses earlier.

Finally, underwriters should take a proactive approach to addressing other future risks. Autonomous AI agents introduce new accountability risks because they may perform high-consequence actions without appropriate human oversight. Underwriters will increasingly need to evaluate this exposure. The paper therefore urges insurers to add questions to their applications asking whether agent actions are logged and attributable to a human owner, whether agents operate with least-privilege access, and whether high-consequence actions require human approval supported by phishing-resistant MFA.

Cyber insurers should not wait for legacy MFA failures to produce a larger and more costly body of claims data before acting. They have already demonstrated that underwriting requirements, pricing incentives, and policyholder guidance can accelerate the adoption of stronger security controls. By making phishing-resistant MFA a core underwriting expectation, insurers can reduce the frequency and severity of identity-driven losses, improve the quality of cyber risk selection, and help organizations prepare for both current AI-enabled phishing threats and the emerging risks created by autonomous AI

systems. We want to thank Yubico for supporting this paper, as well as contributing their knowledge and expertise to the final product.

II. Introduction

Multi-factor authentication (MFA) has long been considered a baseline cybersecurity control, and cyber insurers have played a meaningful role in encouraging its widespread adoption. The threat landscape, however, has changed dramatically. Generative AI has transformed phishing from a relatively crude, error-prone tactic into a fast, scalable, and highly convincing attack method. Attackers are now capable of producing personalized lures, flawless language, and even interactive chatbots designed to manipulate victims in real time. Legacy MFA, which still depends on shareable one-time codes and human judgment calls, is increasingly unable to withstand these advances. Phishing-resistant MFA eliminates shareable authentication factors and prevents users from being tricked into approving fraudulent requests, making it the control best suited to address the current threat environment.

Recent reporting from the U.S. Federal Bureau of Investigation's (FBI) Internet Crime Complaint Center (IC3) highlights the importance of phishing-resistant MFA. In 2025 alone, IC3 received more than 22,000 complaints that referenced AI-related activity, with adjusted losses exceeding \$893 million.⁸ These figures almost certainly understate the true scope of AI's role in cybercrime, since AI involvement is often difficult for victims and investigators to detect or confirm. Even so, the reported losses span nearly every major claims category relevant to cyber insurers, demonstrating that AI-enabled fraud is already an active and growing source of real-world losses.

Despite its effectiveness, phishing-resistant MFA remains underdeployed, creating a growing and urgent challenge for the cyber insurance industry. Phishing already sits upstream of many of the costliest and most frequent claims categories, including business email compromise (BEC), funds transfer fraud (FTF), ransomware, extortion, credential compromise, and data breaches. As AI

⁸https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf

continues to improve the speed and success rate of phishing campaigns, cyber insurers should expect these losses to increase in both frequency and severity, placing new pressure on the financial assumptions that underpin their business model.

Cyber insurance underwriters are uniquely positioned to help close this gap. Since underwriting sits at the intersection of risk assessment, pricing, and insured behavior, even incremental changes to underwriting practices can meaningfully influence how organizations prioritize their security investments.

This paper argues that cyber insurance underwriters do not need to overhaul their practices overnight to make meaningful progress. Instead, underwriters can take incremental steps: asking sharper, more targeted questions about authentication controls, helping insureds strengthen their identity posture during the policy period, and capturing better data on the root causes of claims. Taken together, these steps can help underwriters more accurately segment risk, reduce preventable losses, and steer the market toward controls that reflect the realities of an AI-enabled threat environment.

III. AI-Enabled Phishing and the Changing Threat Landscape

Phishing is a form of social engineering in which attackers trick individuals into disclosing sensitive information, such as passwords or OTPs, by impersonating a legitimate organization or trusted individual. For attackers, persuading a single employee to click a malicious link, approve a login request, or reset an account is often easier and faster than developing the expertise and investing the time needed to bypass security controls directly. Despite its technical simplicity, phishing remains one of the most effective and persistent cyberattack methods, particularly as an initial access vector for malware, ransomware, and account compromise.

The widespread availability of generative AI is fundamentally changing this threat landscape by making phishing attacks faster, cheaper, more convincing, and easier to scale. Since the launch of ChatGPT in

Q4 2022, phishing attacks have reportedly increased by more than 4,100%.⁹ IBM's *Cost of a Data Breach Report 2026* found that more than one in four malicious breaches were AI-enabled, a 56% increase from the previous year, and that these breaches cost an average of \$6 million, roughly \$1 million more than those in which AI use was not reported.¹⁰ In a December 2024 alert,¹¹ the U.S. Federal Bureau of Investigation's (FBI) Internet Crime Complaint Center (IC3) identified several ways attackers are using AI to supercharge phishing campaigns, including:

- ▶ Rapidly generating highly personalized phishing emails, allowing attackers to target more victims with credible, tailored content. IBM's *Cost of a Data Breach Report 2025* found that generative AI reduced the time required to create a convincing phishing email from approximately 16 hours to just five minutes.¹²
- ▶ Producing accurate translations and polished writing that eliminate many of the grammatical errors and awkward phrasing traditionally associated with phishing attempts.
- ▶ Creating convincing content for fraudulent websites and social media profiles, particularly in cryptocurrency and investment scams.
- ▶ Embedding AI-powered chatbots into fraudulent websites to engage victims and encourage them to click malicious links or disclose sensitive information.



The use of these techniques to make attacks more effective is particularly concerning given the fact that phishing already represents one of the largest drivers of losses according to public disclosures from the cyber insurance industry. For example, Coalition reported in its 2026 *Cyber Claims Report* that business email compromise (BEC) and funds transfer fraud (FTF) — both of which often begin with a successful phishing attack — accounted for 58% of claims in 2025. Among FTF claims, 71% involved

⁹<https://www.dashlane.com/blog/genai-fuels-phishing-sophistication>

¹⁰<https://www.ibm.com/forms/mkt-whitepaper-90d12>

¹¹<https://www.ic3.gov/PSA/2024/PSA241203>

¹²<https://www.ibm.com/downloads/documents/us-en/131cf87b20b31c91>

social engineering and 52% originated as BEC; the average loss was \$112,000.¹³ Travelers similarly found that BEC and social engineering consistently ranked among its top three claim categories, representing roughly half of all cyber insurance claims over the past five years.¹⁴ Cowbell likewise identified cybercrime — including phishing, BEC, and FTF — as one of its largest sources of claims and, in its *2026 Claims Report*, cited Anti-Phishing Working Group (APWG) data recording approximately 3.8 million phishing attacks worldwide in 2025.¹⁵ Meanwhile, Resilience connected phishing to higher-severity outcomes, reporting that phishing surged 53% in claim frequency, became the number-one point of failure in 2025, and rose from 21% of incurred losses in 2024 to 50% in 2025.¹⁶ In addition, Resilience attributed a 67% jump in claims tied to phishing, social engineering, and FTF, to the rise of AI tools; such claims rose from 18% in 2024 to 85% in the first half of 2026. “Attackers have used AI to sharpen social engineering across the board, from emails to voice deepfakes, and defensive tactics have to keep pace,” noted the carrier in its *2026 Midyear Claims Report*.¹⁷

As AI makes phishing campaigns more convincing and scalable, insurers are seeing both the frequency and severity of these claims increase. In its *2025 Cyber Risk Report*, Resilience cited research finding that AI-generated phishing campaigns achieved a 54% success rate, compared with 12% for traditional phishing attempts.¹⁸ Chubb similarly identified AI as accelerating the speed and sophistication of cyber threats, highlighting autonomous AI-enabled attacks, privacy risks, ransomware, and supply-chain compromises as key contributors to an increasingly complex cybersecurity risk environment in its *2026 Cyber Claims Report*.¹⁹

¹³<https://www.coalitioninc.com/claims-report/2026>

¹⁴https://info.corvusinsurance.com/hubfs/Ransomware_Reports/Q2_25_Travelers_Cyber_Threat_Report.pdf


¹⁵<https://cowbell.insure/wp-content/uploads/pdfs/CB-US-CyberRoundup-2026ClaimsReport.pdf>

¹⁶<https://cyberresilience.com/blog/2025-cyber-risk-report/>

¹⁷<https://cyberresilience.com/wp-content/uploads/2026/08/2026-Midyear-Cyber-Risk-Report.pdf>

¹⁸<https://cyberresilience.com/blog/2025-cyber-risk-report/>

¹⁹<https://www.chubb.com/us-en/business-insurance/products/cyber-insurance/chubb-cyber-claims-landscape-report.html>




KEY TAKEAWAY FOR UNDERWRITERS

Phishing remains one of the most prevalent and costly attack vectors underpinning a significant share of cyber insurance losses. Generative AI is accelerating this trend by producing phishing campaigns that are more convincing, more personalized, and dramatically easier to scale, increasing both the frequency of successful attacks and the severity of resulting claims.

IV. How Does Phishing-Resistant MFA Reduce an Organization's Cybersecurity Risk and Financial Exposure?

Multi-factor authentication (MFA), also referred to as two-factor authentication or two-step verification, requires a user to present two or more authentication “factors” before gaining access to an account. These authentication factors generally fall into three categories:²⁰



In legacy MFA, a password typically serves as the first factor, while an OTP — delivered via text messages, voice calls, emails, mobile authenticator applications, or hardware tokens — serves as the second. Other legacy MFA implementations rely on push notifications sent to a trusted device, allowing the user to approve or deny a login request with a single tap.

While phishing-resistant MFA may appear similar to legacy MFA to most users, its authentication factors are designed specifically not to be intercepted, shared, or reused by an attacker. Phishing-resistant MFA always has a possession factor that controls and protects authentication credentials. The possession factor may take the form of a security key (e.g., a physical hardware token), a platform authenticator (e.g., an authenticator embedded directly into a user's smartphone), or a synced software authenticator (e.g., a password manager) that the user controls. The authenticator will respond only to a pre-registered domain or website. This design not only protects against phishing but also other common attack vectors such as adversary-in-the-middle (AiTM) attacks, push bombing (also known as push fatigue), and SIM swapping, all of which attackers use to gain unauthorized access to accounts. In the case of a security key, the user is also required to tap or touch the key, showing intent to authenticate versus a remote hacker trying to log in. The attack surface is significantly reduced to

²⁰https://csrc.nist.gov/glossary/term/multi_factor_authentication

who has access to the authenticator and the ability to unlock it. *(For more information on how to distinguish between different types of MFA solutions, review the Technical Annex.)*

Thanks to these protections, phishing-resistant MFA can both reduce the likelihood of a successful cyberattack and mitigate the severity of an attack that does occur. This reduces the financial exposure of not only the organization that might otherwise suffer an incident, but also to their cyber insurers, whose financial exposure depends on claims.

Phishing-resistant MFA reduces the likelihood of an incident in several ways, including by:

- ▶ **Preventing stolen credentials from being used for account takeover.** When a password is paired with a phishing-resistant factor, such as a FIDO security key, stealing the password alone is insufficient to access the account. Even if an attacker directs a user to a fraudulent login page and captures the password, the authenticator will reject the request because it is not associated with the legitimate service domain.
- ▶ **Blocking AiTM phishing attacks.** AiTM phishing kits attempt to relay login activity in real time and capture credentials, one-time passcodes, or session tokens. Phishing-resistant MFA disrupts this process because the authenticator will not complete a request initiated through a domain other than its registered domain.
- ▶ **Reducing reliance on user judgment during authentication.** Security awareness is still important, but with phishing-resistant MFA users no longer have to determine whether a login page is legitimate. The FIDO2 passkey protocol performs that verification as part of the authentication process.

Initially protecting key accounts that have access to sensitive systems or information with phishing-resistant MFA can reduce the severity of a cybersecurity incident. The consequences of a compromised account vary depending on the user's responsibilities and privileges. For example, gaining access to an email account can enable BEC and FTF; gaining access to a cloud administrator account can create opportunities for ransomware, data theft, or broader infrastructure compromise; and gaining access to a help-desk account may allow an attacker to reset other users' credentials and move deeper into the organization. Protecting these high-risk accounts can prevent an initial intrusion from developing into a major operational or financial event. Once key accounts are protected, it is imperative to continue to

protect all accounts. With the advancement of AI, there is a greater chance that non-privileged accounts can be an entry point to gain elevated privilege.

This reduction in risk can produce a meaningful financial return for organizations that implement phishing-resistant MFA. For example, Forrester Consulting's January 2026 *Total Economic Impact of Yubico YubiKeys* study found that a composite enterprise that replaced earlier forms of MFA and OTPs with YubiKeys achieved a 265% return on investment and a net present value of \$5.3 million over three years.²¹ The study also estimated that the deployment reduced exposure to breach costs from attacks addressed by the technology by 99.9%.²² For organizations with cyber insurance, a reduction in breach-related losses may also reduce insurers' claims exposure, allowing those savings to extend beyond the insured organization.



265%

return on investment
from FIDO2 hardware
security keys

Forrester / Yubico, Jan. 2026



\$5.3M

net present value over
three years for a
composite enterprise

Forrester / Yubico, Jan. 2026



99.9%

reduction in exposure to
breach costs from
addressed attacks

Forrester / Yubico, Jan. 2026

KEY TAKEAWAY FOR UNDERWRITERS

Phishing-resistant MFA reduces both the likelihood and severity of an incident by eliminating the shareable codes and human judgment calls that legacy MFA still depends on. Organizations deploying phishing-resistant MFA have reported dramatic reductions in breach exposure alongside strong returns on investment — savings that can also reduce insurers' downstream claims exposure.

²¹<https://tei.forrester.com/go/yubico/yubikeys/?lang=en-us>

²²<https://www.yubico.com/press-releases/yubico-roi/>

V. A Cyber Insurance Underwriter's Perspective on Responding to AI-Enabled Phishing

As cyber insurance underwriters face increased financial exposure due to AI-enabled phishing threats, they must determine whether their key assumptions and practices have evolved enough to account for the modern threat environment, and thus properly evaluate the effectiveness of a prospective insured's cybersecurity controls. To better understand how cyber insurance underwriters are currently approaching this issue, the Center for Cybersecurity Policy and Law ("the Center") conducted interviews with industry experts, reviewed publicly available data, and undertook a targeted literature review. The key insights and themes that emerged from this research are reflected below.

Are Underwriters Aware of the Risks Posed by AI-Enabled Phishing?

As noted in Section III, public claims data clearly demonstrates that phishing is a major loss driver for insureds. Cyber insurance experts also routinely described AI as an accelerant that improves the quality, speed, and scale of phishing, which is consistent with FBI IC3 findings in its December 2024 alert.²³ However, interviewees noted that claims files often do not capture enough technical detail to prove an attacker leveraged generative AI in a specific incident. In practice, claims are then coded as BEC, social engineering, credential misuse, ransomware, or FTF, rather than being separately or additionally identified as "AI-enabled phishing." Therefore, even if underwriters recognize the risks posed by AI-enabled phishing, the lack of data collection — and of consistent forensic and claims taxonomies for the data that is collected — makes it difficult to actually measure and price risk. This evidence gap weakens the business case for updating underwriting practices, despite broad recognition that the threat is growing.

²³<https://www.ic3.gov/PSA/2024/PSA241203>

Are Underwriters Aware of the Differences in Effectiveness Between Legacy and Phishing-Resistant MFA?

Compared with the large body of claims reporting on ransomware, BEC, and FTF, the Center found that there are relatively few public, insurance-specific analyses that quantify the loss difference in effectiveness between legacy MFA and phishing-resistant MFA. However, there is evidence to suggest that awareness of the importance of phishing-resistant MFA is growing in certain pockets of the cyber insurance underwriting community.

For example, Marsh McLennan Cyber Risk Intelligence Center's August 2025 report, titled *Cybersecurity signals: Connecting controls and incident outcomes*, found that phishing-resistant MFA is associated with an 18% reduced likelihood of a breach, compared with 8% from legacy MFA.²⁴ Marsh McLennan's next iteration of its *Cyber Self-Assessment* — a digital tool used for applying for cyber insurance — plans to draw more attention to this issue by providing an organization with resources to evaluate the maturity of its identity program in greater depth, including its resilience against increasingly sophisticated phishing techniques.

Similarly, Travelers' Q2 2025 *Cyber Threat Report* provides a practical illustration of why the distinction between different types of MFA matters. In the report, Travelers cites a specific occurrence in which a manufacturing firm required MFA for all accounts, but allowed SMS authentication as the second factor. An attacker exploited that weakness using an AiTM attack, captured credentials, compromised the CFO's account, and used that access to facilitate fraudulent wire transfers totaling £150,000. Travelers later identified both the weak MFA configuration and the lack of out-of-band payment verification as key control failures. Its guidance recommends MFA for email and financial systems, ideally using phishing-resistant MFA such as hardware tokens or app-based authenticators, and separately recommends phishing-resistant MFA for all remote access and email.²⁵

²⁴<https://www.marsh.com/en/services/cyber-risk/insights/cybersecurity-signals.html>

²⁵https://asset.trvstatic.com/download/assets/Q2_25_Travelers_Cyber_Threat_Report.pdf/a86959b888c011f0a6e482ee1c9c7491

The Center's interviews with cyber insurance experts also reinforce the idea that underwriters increasingly acknowledge the importance of differentiating between legacy and phishing-resistant MFA. Many interviewees acknowledged that *"Has this organization implemented MFA?"* is no longer an adequate underwriting question since a prospective insured can technically answer "yes" even when its MFA deployment is outdated, incomplete, or easily bypassed. This may include legacy MFA that is only partially deployed, is not enforced for privileged users, allows SMS or email-based fallback, can be bypassed through help-desk recovery, or is not applied to vendors and contractors.

How Are Underwriters Assessing Phishing-Related Cybersecurity Risk Before Binding Coverage?

Since most cyber insurance applications are proprietary, it is difficult to determine precisely how often underwriters ask prospective insureds about phishing-resistant MFA or other identity controls. It is even more difficult to assess how these controls are weighed relative to other underwriting factors, as underwriters generally do not disclose their application logic, scoring methodologies, or the criteria they consider material when segmenting risk and pricing policies.

That said, cyber insurance experts interviewed by the Center consistently noted that, while phishing-resistant MFA is increasingly recognized as an important indicator of resilience, incorporating detailed questions about identity security into the underwriting process remains constrained by market realities. Brokers often process large volumes of applications, applicants expect a streamlined experience, and many small and medium-sized enterprises (SMEs) may lack the technical expertise to answer detailed questions about identity architecture. In a competitive market, underwriters may also be reluctant to introduce additional friction into the application process. As a result, underwriters may rely on a smaller set of high-level questions, narrative responses, or applicant self-attestation rather than detailed questions to validate whether identity controls like phishing-resistant MFA have been implemented effectively. Beginning on September 1, 2026, Microsoft is discontinuing native support

for SMS and rolling out passkeys as the default authentication experience for Microsoft Entra ID users.²⁶

KEY TAKEAWAY FOR UNDERWRITERS

Underwriters increasingly recognize that phishing-resistant MFA offers materially stronger protection than legacy MFA, but limited claims data, competitive pressure to streamline applications, and reliance on applicant self-attestation have made it difficult to translate that awareness into more probing underwriting questions. Simply asking “*Has this organization implemented MFA?*” without also asking whether that MFA is phishing-resistant, and probing further into what the deployment looks like in practice, does not give underwriters an accurate understanding of the risks an organization faces before binding coverage.

VI. Current State of Adoption of Phishing-Resistant MFA

Due to its effectiveness, phishing-resistant MFA has emerged as one of the most widely recommended cybersecurity controls for both public- and private-sector organizations. It is explicitly recommended in guidance published by several national cybersecurity agencies, including:

- ▶ U.S. Cybersecurity and Infrastructure Security Agency's (CISA) *Secure Your Business* guidance²⁷
- ▶ UK National Cyber Security Centre's (NCSC) *Multi-factor Authentication for Your Corporate Online Services* guidance²⁸
- ▶ European Union Agency for Cybersecurity's (ENISA) *Technical Implementation Guidance for the NIS 2 Directive (Directive (EU) 2022/2555)*²⁹

²⁶<https://www.microsoft.com/en-us/security/blog/2026/07/13/microsoft-entra-id-security-updates-passkeys-are-the-default-authentication-method-in-entra-id/>


²⁷<https://www.cisa.gov/audiences/small-and-medium-businesses/secure-your-business>

²⁸<https://www.ncsc.gov.uk/collection/mfa-for-your-corporate-online-services>

²⁹<https://www.enisa.europa.eu/publications/nis2-technical-implementation-guidance>

- ▶ Singapore Cyber Security Agency's (CSA) *Advisory on Defending Against Multi-Factor Authentication (MFA) Bypass Attacks*³⁰
- ▶ Australian Signals Directorate's (ASD) *Essential Eight*³¹
- ▶ Japan Financial Services Agency's (FSA) public awareness campaign on phishing-resistant MFA³²

By most accounts, adoption of phishing-resistant MFA is increasing. The FIDO Alliance's *State of Passkeys 2026: Global Consumer and Workforce Report* indicates that passkeys — a leading form of phishing-resistant MFA — have reached significant scale, with more than five billion passkeys in active use globally. (For more information about passkeys, review the *Technical Annex*). Across ten surveyed markets, including the United States, United Kingdom, France, Germany, Australia, Singapore, Japan, South Korea, China, and India, 90% of consumers reported being familiar with passkeys, while 75% had enabled them for at least some accounts. Adoption within organizations is also approaching the mainstream, with 68% reporting that they have deployed, are piloting, or are rolling out passkeys for employee authentication.³³ Similarly, Yubico reports that more than 8,000 enterprise customers have deployed its YubiKey security keys to strengthen workforce authentication. However, other accounts suggest adoption, while growing, still remains low. According to Okta's *Secure Sign-In Trends Report 2025*, approximately 70% of organizations have deployed some form of MFA, but only 14% have implemented phishing-resistant MFA.³⁴



³⁰<https://www.csa.gov.sg/alerts-and-advisories/advisories/ad-2024-020/>

³¹<https://www.cyber.gov.au/business-government/asds-cyber-security-frameworks/essential-eight>

³²<https://www.fsa.go.jp/news/r7/sonota/20260416-2/20260416-2.html>

³³<https://fidoalliance.org/the-state-of-passkeys-2026-global-consumer-and-workforce-report/>

³⁴<https://www.okta.com/newsroom/articles/secure-sign-in-trends-report-2025/>

INDUSTRY EXAMPLE — T-MOBILE US AND YUBICO

In January 2025, Yubico, a security technology company, announced it had worked with T-Mobile US, a major U.S. wireless carrier and telecommunications provider, to deploy FIDO2 YubiKey security keys to employees companywide as well as to vendors and authorized retail partners. The rollout, completed in late 2023, replaced legacy MFA methods that were more vulnerable to phishing. According to T-Mobile Chief Security Officer Jeff Simon, the company completed the deployment in less than three months after receiving the keys.³⁵

The deployment highlights two key points that are relevant to cyber insurance underwriters:

- ▶ Large-scale phishing-resistant MFA is operationally feasible. T-Mobile extended the rollout beyond employees to vendors and authorized retail partners. This is especially significant given third-party access often remains a weak point even in otherwise mature identity programs.
- ▶ A rapid rollout is possible when deployment is treated as an enterprise program. A less-than-three-month rollout across a complex organization suggests that implementation barriers can be managed when leadership, identity architecture, logistics, support, and communications are aligned.

Regardless, adoption varies considerably across different types of organizations. Cloud-native organizations often find it relatively straightforward to deploy passkeys and other phishing-resistant authentication methods through the identity platforms they already use. Similarly, highly regulated enterprises frequently have mature identity governance programs, centralized identity providers, and dedicated cybersecurity budgets that make implementation more practical. By contrast, many SMEs continue to rely on SMS-based authentication, email OTPs, basic push notifications, or only partial MFA coverage because these approaches are widely marketed as effective security controls, and/or sufficient to satisfy existing compliance requirements. However, despite these challenges, SMEs may eventually adopt phishing-resistant MFA at higher rates than other types of organizations because many rely on passkey-capable, cloud-based work environments, such as Google Workspace and Microsoft 365, that support streamlined deployment.


³⁵<https://www.yubico.com/press-releases/yubico-and-t-mobile-deployment-of-200000-phishing-resistant-yubikeys-enhances-un-carriers-work-systems-security/>

VII. Best Practices for Deploying Phishing-Resistant MFA

To maximize its security benefits, organizations should deploy phishing-resistant MFA as part of a broader identity and access management strategy rather than as a standalone control. While phishing-resistant MFA on its own significantly reduces the risk of credential theft and account compromise, it is most effective when combined with complementary security measures such as least-privilege access controls, privileged access management, device security, user awareness training, continuous monitoring, and incident response capabilities.

Organizations should also carefully review account recovery and fallback authentication mechanisms. If users can bypass phishing-resistant MFA through a weaker, phishable authentication method, the account may remain vulnerable despite the deployment of stronger authentication. As a result, organizations should seek to minimize or eliminate fallback mechanisms that undermine the security benefits of phishing-resistant MFA.

Where organization-wide deployment is not immediately feasible, phishing-resistant MFA should be prioritized for the most privileged and high-risk accounts, including administrators, executives, finance personnel, developers, and others with access to sensitive systems or data. These accounts are often the primary targets of sophisticated attackers and therefore provide the greatest risk reductions when adequately protected. A pragmatic deployment strategy should start with the accounts and systems most likely to cause material loss:



- Domain administrators

- ▶ Cloud administrators
- ▶ Email administrators
- ▶ Executives
- ▶ Finance and payment personnel
- ▶ IT help-desk users
- ▶ Developers with production or code-signing access
- ▶ Users with access to sensitive data
- ▶ Users with access to remote-access infrastructure
- ▶ Vendors and contractors with privileged or persistent access

Prioritizing these accounts, however, should be viewed as an initial deployment phase rather than the end state. In addition to making phishing more advanced, frontier AI is improving attackers' abilities to identify software vulnerabilities and target weaknesses across an organization's environment. In this context, attackers may not need to compromise a highly privileged account directly. A lower-privileged user account can provide an initial foothold from which an attacker may seek to escalate privileges, move laterally, and ultimately gain access to more sensitive systems and data. For this reason, organizations should work toward extending phishing-resistant MFA to the broader workforce over the medium to long term. Protecting lower-privileged accounts reduces the number of potential entry points available to attackers and strengthens the organization's overall resistance to credential-based compromise.

Finally, organizations with limited cybersecurity resources or highly complex IT environments should consider leveraging systems integrators, managed security providers, or other specialized implementation partners to support deployment efforts. These organizations can help identify critical accounts, integrate phishing-resistant MFA into existing identity infrastructure, manage technical challenges, and ensure that implementations are configured correctly and consistently across the enterprise.

KEY TAKEAWAY FOR UNDERWRITERS

To be most effective, all users need to use phishing-resistant MFA, but initial deployments should prioritize the most privileged and high-risk accounts, eliminate weak fallback authentication methods, and integrate the control into a broader identity and access management program.

VIII. Recommendations for Underwriters

As established in Section IV, phishing-resistant MFA has the potential to reduce risk for insureds and, in turn, limit insurers' financial exposure. However, this paper also argues that adoption of these solutions continues to lag behind where it should be. Fortunately, cyber insurers are well positioned to accelerate implementation because they operate at the intersection of risk assessment, pricing, claims, and digital resilience.

The recommendations below explain how underwriters can encourage broader adoption by improving the questions included in cyber insurance applications, helping insureds strengthen their controls during the policy period, and capturing better data on the underlying causes of claims.



Modernize Application Questions

Replace low-signal questions with specific, high-signal ones on identity controls



Strengthen Identity Posture

Use the policy period to help insureds improve authentication maturity




Capture Better Claims Data

Build a consistent taxonomy for the identity-control failure behind a loss

Improve Questions for Cyber Insurance Applications

The following table proposes a prioritized set of high-signal, practical underwriting questions that underwriters can use to evaluate an organization's MFA implementation and identity controls. The questions are designed to be more specific than simply asking “*Has this organization implemented MFA?*” and to serve as proxies for broader maturity of its cybersecurity program. They have been prioritized into three tiers representing company expectations from lowest to highest maturity to help

underwriters align application depth to the size of the entity they are insuring. Tier 1 would be applicable for SMEs with limited budget, staffing, and revenue at risk. Tier 2 would be applicable for mid-sized organizations with a regular security program, but still facing limitations on staffing and budget. Tier 3 would be applicable to larger enterprises with a mature security program.



Tier 1 Questions — SMEs

Tier	Underwriting Objective	Updated Application Question	How This Improves Risk Assessment
1	Identify MFA scope	Does the organization require MFA for email, remote access, cloud applications, and privileged accounts?	Determines whether MFA covers common claim pathways, not just a narrow subset of access.
1	Distinguish MFA method	Which authentication methods are permitted, including SMS, email one-time passwords, voice calls, push notifications, time-based one-time passwords, passkeys, FIDO2/WebAuthn security keys, certificate-based authentication, or other methods?	Reveals whether the control is phishable, relayable, or phishing-resistant.
1	Identify phishing-resistant coverage	Is phishing-resistant MFA required for administrators, executives, finance personnel, IT and help-desk users, developers with production access, and other high-risk users?	Protects accounts most likely to create high-severity loss.
1	Assess resilience of payment workflows	Are payment changes and other high-risk financial requests verified through an out-of-band channel and subject to dual approval?	BEC and social engineering can succeed even when login controls work.

Tier 2 Questions — Mid-Market

Tier	Underwriting Objective	Updated Application Question	How This Improves Risk Assessment
2	Test fallback resilience	Are weaker fallback methods, such as SMS, email one-time passwords, voice calls, or security questions allowed for accounts otherwise protected by stronger authentication?	Weak fallback methods can undermine phishing-resistant deployment.
2	Evaluate account recovery	Can users bypass MFA through help-desk resets, account recovery, emergency access procedures, or legacy authentication protocols?	Many social-engineering attacks target process bypass rather than the MFA protocol.
2	Evaluate vendor and contractor risk	Are vendors, contractors, and third-party service providers subject to the same authentication requirements as employees?	Third-party access can create a back door around employee controls.
2	Assess privileged access governance	Does the organization maintain separate administrative accounts for privileged activities and/or only provide privileged permissions when necessary?	Indicates whether the organization understands privilege separation and blast-radius reduction.
2	Assess identity lifecycle	How are privileged accounts issued, reviewed, suspended, and revoked when employees or contractors change roles or leave the organization?	Dormant and stale accounts increase compromise risk.

Tier 3 Questions — Enterprise

Tier	Underwriting Objective	Updated Application Question	How This Improves Risk Assessment
3	Assess enforcement	Is phishing-resistant MFA enforced through conditional-access policies, equivalent technical controls, or identity-provider configuration?	Policies must be technically enforced, not only documented.
3	Assess monitoring	Does the organization monitor for anomalous logins, authenticator changes, impossible travel, token theft, Open Authorization (OAuth) abuse, suspicious account recovery, and dormant privileged accounts?	Detection and response matter when attackers obtain valid credentials or sessions.
3	Assess high-risk change events	Are identity controls reassessed after mergers, cloud migrations, identity-provider changes, major incidents, or the adoption of AI agents?	Identity risk changes as the business changes.

Help Insureds Improve During the Policy Period

Underwriting should not be viewed as the only point at which insurers can influence an insured's cybersecurity posture. Insurers can also use the policy period to encourage stronger identity controls, support implementation, and create incentives for measurable improvements. For example, underwriters could:

- ▶ Offer broker-facing education on the difference between legacy and phishing-resistant MFA
- ▶ Provide insured-facing implementation guidance for common identity environments
- ▶ Partner with systems integrators, managed security service providers, identity specialists, and managed detection and response providers
- ▶ Offer premium credits, retention benefits, or improved terms for verified phishing-resistant MFA deployment
- ▶ Support hardware security key pilot programs for privileged users
- ▶ Build tabletop exercises that include BEC, MFA bypass, account recovery, and deepfake scenarios

- ▶ Integrate phishing-resistant MFA guidance into post-claim remediation plans

Capture Data on the Underlying Causes of Cyber Insurance Claims

As suggested by numerous interviewees for this paper, the cyber insurance industry needs better data on the underlying causes of claims. Many claims are categorized by outcome, such as BEC, FTF, ransomware, data breach, or extortion, rather than by the identity control failure that enabled the loss.

To address this, cyber insurers should develop a more consistent taxonomy for identity-related claims. Specifically, forensic vendors, breach counsel, claims teams, and risk engineers should record whether an incident involved:

- ▶ No MFA
- ▶ Partial MFA
- ▶ Phishable MFA
- ▶ SMS or email OTP
- ▶ Push bombing (push fatigue)
- ▶ AiTM phishing (fake website)
- ▶ Session hijacking
- ▶ OAuth consent abuse
- ▶ Compromised vendor credentials
- ▶ Account-recovery bypass
- ▶ Help-desk social engineering
- ▶ Lack of phishing-resistant MFA

IX. Future-Proofing for Agentic AI

While AI-enabled phishing poses an immediate threat that cyber insurers will need to account for in their underwriting practices, it is not the only challenge they will face as AI and other emerging

technologies reshape the digital landscape. Another significant risk on the horizon stems from AI agents, software systems that use AI to perceive context, reason through objectives, and execute multi-step workflows with limited human oversight. Organizations are increasingly deploying AI agents across their digital environments to write and deploy code, initiate financial workflows, access legal and compliance materials, and orchestrate security operations.


Deployment of AI agents in organizations raises a range of governance and accountability questions. Chief among them is how organizations should manage autonomous decision-making and determine accountability when an AI agent takes an incorrect, high-consequence action without explicit human authorization. Consider a scenario in which an AI agent executes a multimillion-dollar financial transaction to the wrong bank account. Although the agent possesses technically valid credentials to complete the transaction, no human directed or authorized the payment, and the funds cannot be recovered. In such circumstances, determining who bears responsibility for the loss becomes far more complex than in other types of incidents.

Until these types of controls become commonplace and governance practices for AI agents mature, cyber insurers will likely have to contend with claims arising from these types of incidents. Expanding underwriting practices to account for phishing-resistant MFA and other modern identity controls — as recommended in this paper — is an important first step toward a broader identity resilience underwriting model that is better equipped for the age of AI agents. Beyond addressing AI-enabled phishing, these questions can help underwriters assess whether an organization has the identity governance, access controls, and operational maturity necessary to deploy AI agents responsibly and manage the new risks they introduce.

For insurers seeking to take a more proactive approach to AI agent risk, underwriting applications could begin incorporating basic AI identity governance questions for organizations that deploy AI agents or AI-enabled workflows. For example:

- ▶ Does the organization maintain an inventory of AI tools and agents?
- ▶ Are AI agents assigned least-privilege access?

- ▶ Are agent actions logged and attributable to a human owner?
- ▶ Can agents access sensitive data, payment systems, production systems, code repositories, or customer environments?
- ▶ Are Application Programming Interface (API) keys, service accounts, delegated access to applications (OAuth), and machine identities inventoried and regularly rotated?
- ▶ Does the organization have an AI acceptable-use policy and a formal approval process for new AI tools?
- ▶ Are high-consequence agent actions subject to human approval using phishing-resistant authentication methods? Are the approvals appropriately recorded?



Similar to the questions recommended in Section VIII, these questions reflect a broader shift in cyber insurance underwriting. Rather than simply asking whether an organization has security controls in place, underwriters should evaluate whether those controls are deployed in a manner that meaningfully reduces risk.

KEY TAKEAWAY FOR UNDERWRITERS

Expanding underwriting practices to account for phishing-resistant MFA and other modern identity controls is an important first step, not only to address AI-enabled phishing, but also to move toward a broader identity resilience underwriting model better equipped for the age of AI agents. Underwriters who want to be more proactive about future risks should also consider updating applications to ask specific questions about AI agents or AI-enabled workflows.

X. Conclusion

It is becoming increasingly apparent that legacy MFA while once considered good enough is no longer good enough in an AI world. As attackers increasingly harness generative AI to craft phishing campaigns that are more convincing, more personalized, more scalable, and harder for employees to detect, only phishing-resistant MFA is properly equipped to address the current threat.

However, phishing-resistant MFA deployment continues to lag. This poses an acute problem for the cyber insurance industry, since AI-enabled phishing is a key access vector behind many of the most costly and frequent claims categories, including business email compromise, funds transfer fraud, ransomware, extortion, credential compromise, and data breach. Current trends give every reason to expect these claims to grow in both frequency and severity.

Fortunately, cyber insurance underwriters have a real opportunity to leverage their reach and influence to help drive adoption of phishing-resistant MFA among insureds. Underwriters are increasingly aware of the risks posed by AI-enabled phishing and of the protective value phishing-resistant MFA provides. The key challenge in boosting adoption is that cyber insurance underwriting practices — as is the case in the broader insurance industry — are built on historical claims data and tend to evolve only after a trend has generated enough losses to become clearly visible. That lag was tolerable when threats evolved slowly; it is far less tenable now. Indeed, the acceleration of attacker capabilities and the potential for concentrated, correlated losses threaten to strain the financial assumptions underlying the cyber insurance business model. Cyber insurance practices must evolve in step.

Recognizing that transforming underwriting practice overnight is unrealistic, this paper argues for incremental steps cyber insurance underwriters can take to drive broader awareness and adoption of phishing-resistant MFA. Insurers do not need to turn every application into a full IAM audit, but they can ask sharper, more targeted questions about authentication controls. The answers to those questions can help underwriters segment risk more accurately, reduce preventable claims, and steer policyholders toward controls that match today's threat environment. Beyond underwriting, insurers

are also well positioned to help insureds strengthen their authentication posture during the policy period, and to capture better data on the root causes of claims. That data can, in turn, sharpen future underwriting and help the industry keep pace with attacker innovation rather than trail behind it.

Taken together, these steps can help ensure the cyber insurance model stays financially sound while improving the overarching resilience and security of organizations in a fast-changing digital environment.

Technical Annex. Legacy vs. Phishing-Resistant MFA

This annex is intended to help underwriters and risk engineers ask better questions on cyber insurance applications. As discussed throughout this paper, underwriters should move beyond simply asking whether an applicant has deployed multi-factor authentication (MFA) and instead determine whether it has implemented phishing-resistant MFA.

To do so, both underwriters and their risk engineering teams need a practical understanding of how phishing-resistant MFA differs from legacy MFA. Although the distinction between MFA and phishing-resistant MFA wasn't carefully considered in the past, with regulations now increasingly mandating phishing-resistant MFA.

This annex provides the technical foundation underwriters and risk engineers need to ask informed questions, interpret applicants' responses, and distinguish between authentication methods that offer fundamentally different levels of security.

What is Legacy MFA?

Legacy MFA encompasses a wide range of authentication methods. Most commonly, an alphanumeric password with specific complexity requirements — such as a minimum length and a combination of uppercase and lowercase letters, numbers, and special characters — serves as the first factor, and a one-time passcode (OTP) serves as the second. These OTPs may be delivered in several different forms, such as a text message, voice call, email, mobile authenticator application, or hardware token that displays a rotating code. Other legacy MFA implementations rely on push notifications sent to a trusted device, allowing the user to approve or deny a login request with a single tap.

Although these solutions differ in their implementation, they share a common characteristic: they require users to receive, view, enter, or approve an authentication factor that can ultimately be intercepted, relayed, manipulated, or obtained through social engineering. Legacy MFA outputs are

not bound to a registered device and as a result the login system does not know if the code was sent by an attacker. Common attack techniques include:³⁶

- ▶ **Phishing** — An attacker uses social engineering techniques — such as fraudulent emails, fraudulent websites, or impersonation of a trusted individual — to persuade a user to provide an MFA code or otherwise authorize the attacker's attempt to access an account.
- ▶ **Push Bombing (Push Fatigue)** — An attacker repeatedly sends authentication approval requests to a user's device, prompting the user to eventually approve a request accidentally or out of frustration in order to stop the notifications.
- ▶ **SS7 Protocol Exploitation** — An attacker exploits vulnerabilities in the Signaling System No. 7 (SS7) telecommunications protocol to intercept MFA codes delivered through text messages or voice calls.
- ▶ **SIM Swapping** — An attacker uses social engineering or other fraudulent means to convince a mobile carrier to transfer a victim's phone number to a SIM card under the attacker's control, allowing MFA codes delivered through text messages or voice calls to be received by the attacker.
- ▶ **Adversary-in-the-Middle (AiTM)** — An attacker positions themselves between a user and a legitimate service, relaying communications between both parties while capturing credentials, authentication codes, and session tokens.

³⁶<https://www.cisa.gov/sites/default/files/publications/fact-sheet-implementing-phishing-resistant-mfa-508c.pdf>

The table below provides a more detailed breakdown of which legacy MFA solutions are vulnerable to each of these attack techniques.

Category	MFA Authenticator	Overview	Threat
One-Time Passcode (OTP)	Text Message (SMS)	User receives an OTP via text message.	Vulnerable to: Phishing, SS7, SIM Swap, AiTM N/A attacks: Push bombing
	Voice Message	User receives a phone call that gives them an OTP.	Vulnerable to: Phishing, SS7, SIM Swap N/A attacks: Push bombing
	Email	User receives an OTP via email.	Vulnerable to: Phishing N/A attacks: Push bombing, SS7, SIM Swap
	OTP Hardware Token	User obtains an OTP from a dedicated hardware device that generates rotating authentication codes.	Vulnerable to: Phishing N/A attacks: Push bombing, SS7, SIM Swap
Push-based Authenticator Apps	Mobile push with number matching	User receives a push notification and approves the login request by entering a 2-digit number displayed on the screen into the app.	Vulnerable to: Phishing Resistant to: Push bombing N/A attacks: SS7, SIM Swap
	Mobile push without number matching	User receives a push notification and approves or denies the login request with a single tap.	Vulnerable to: Phishing, Push bombing N/A attacks: SS7, SIM Swap

Note: While this table covers the most common legacy MFA solutions used in enterprise environments, it is not exhaustive. Other authentication methods, such as email-based magic links, security questions,

and other forms of knowledge- or possession-based authentication, are also used in certain circumstances.

What is Phishing-Resistant MFA?

Although phishing-resistant MFA solutions may appear similar to legacy MFA solutions from the perspective of the user, there are fundamental differences between the two. Whereas legacy MFA relies on authentication factors that can be shared, intercepted, or relayed to another party, phishing-resistant MFA is designed specifically so that the factor cannot be intercepted, shared or reused by an attacker. As a general rule, **if the outputs of an authentication factor can be disclosed to another person that is not present and used to grant access to their device, it is not phishing-resistant.**

Most phishing-resistant MFA solutions leverage public-key cryptography to avoid the need for OTPs and other shareable credentials. When a user enrolls in a phishing-resistant MFA solution, it generates a unique cryptographic key pair — a public key and a private key. The private key is securely stored within and bound to an authenticator, such as a user's device, a hardware security key like a YubiKey, or a password manager that holds users' credentials. The corresponding public key is registered with the online account or service being accessed, such as a bank, email provider, or workplace application.

When the user attempts to log into their account, the service sends a request to the authenticator asking it to prove possession of the private key. The user's authenticator uses the private key to generate a response, which the service verifies using the corresponding public key it has on file. **Since the private key never leaves the user's authenticator, there is no code, password, or credential for the user to read, copy, share, or enter into a website.** Additionally, the authenticator will only respond to pre-registered sites, blocking requests from fraudulent locations. Therefore, attackers cannot capture and reuse the authentication factor in the same way they can with legacy MFA methods. This combination of non-shareable credentials and website verification is what makes phishing-resistant MFA effective against phishing, push fatigue, SS7 Protocol Exploitation, SIM swapping, AiTM, and other techniques that routinely defeat legacy MFA solutions.

FIDO2 Authentication and Passkeys

FIDO2 is the most widely available and widely deployed form of phishing-resistant MFA. Developed by the FIDO Alliance and the World Wide Web Consortium (W3C), it provides a standardized framework for authenticating users through passkeys. FIDO2 consists of two main components:

- ▶ WebAuthn (Web Authentication) — Defines how websites and applications create and use credentials through web browsers.
- ▶ CTAP2 (Client to Authenticator Protocol 2) — Defines how the client (browsers and/or operating systems) communicate with authenticators such as security keys, smartphones, and laptops.

FIDO2 evolved from an earlier standard known as FIDO Universal 2nd Factor (U2F). FIDO2 authenticators support user verification through PINs and biometrics, such as fingerprints and facial recognition. Whereas U2F functioned only as a second authentication factor, modern FIDO2 implementations can be used either as a second factor or as the sole authentication factor. When used as the sole factor, the credential is referred to as a “passkey.” There are two main distinctions of passkeys based on how they are stored:

- ▶ Device-bound passkeys — Passkeys that are stored locally on a specific piece of hardware, such as a security key, and cannot be synced to the cloud or copied to other devices. They offer the highest level of security and auditability. NIST has stated that device-bound passkeys meet the requirements for its highest authentication level, Authenticator Assurance Level 3 (AAL3). They are most typically used in enterprise use cases.
- ▶ Synced passkeys — Passkeys that are synchronized across multiple devices of the same operating system (e.g., Apple or Android), or through software like password managers. NIST has stated that synced passkeys can meet Authenticator Assurance Level 2 (AAL2). They are most typically used in consumer-facing use cases, and offer reduced auditability compared to device-bound passkeys.

In a FIDO2 deployment, passkeys can be stored in several different types of authenticators. While all of these authenticators are generally considered phishing-resistant, they provide different levels of security, convenience, and assurance.

- ▶ Security Keys — Standalone physical hardware tokens that store device-bound passkeys and connect to computers or mobile devices through USB, NFC, or BLE. Generally considered the most secure option.
- ▶ Platform Authenticators — Authenticators embedded directly into a user's laptop, smartphone, or tablet, using dedicated security hardware such as a TPM, TEE, or Apple's Secure Enclave. Provide a high level of security with the convenience of syncing across devices.
- ▶ Software Authenticators — Software-based authenticators and password managers that store and synchronize passkeys across multiple devices. Security is partially dependent on the security of the password manager account itself.

The FIDO Alliance maintains a metadata service so organizations can validate certified devices that can be used to make authentication decisions.

PKI-Based Authentication

While FIDO2 is the most widely deployed form of phishing-resistant MFA, governments and enterprises also rely on authentication methods based on Public Key Infrastructure (PKI). Unlike FIDO2, which is broadly supported across modern websites and applications and is designed to be decentralized and relatively frictionless for users, PKI-based authentication relies on cryptographic keys that are issued, managed, and trusted through a centralized Certificate Authority (CA). In some cases, these credentials are contained within a security chip embedded in a smart card, and the card must be connected to a device for the user to authenticate using the correct password or PIN.

Examples of PKI-based MFA include the U.S. government's Personal Identity Verification (PIV) card and Common Access Card (CAC) standard and Microsoft's Certificate-Based Authentication (CBA). Some modern hardware authenticators such as security keys are capable of storing both FIDO2/WebAuthn and PKI-based credentials.

The table below provides a breakdown of different phishing-resistant MFA solutions and examples of commonly deployed products and authenticators used by enterprise organizations:

Category	MFA Authenticator	Overview
FIDO2 Authenticators	Security Keys	Separate physical devices that store passkeys and connect to computers or mobile devices through USB, Near Field Communication (NFC), or Bluetooth Low Energy (BLE).
	Platform Authenticators	Authenticators embedded directly into a user's laptop, smartphone, or tablet. Apple and Google versions will sync passkeys across devices of the same operating system.
	Synced Software Authenticators	Software-based authenticators and password managers that store and synchronize passkeys across multiple devices. Some authenticators support device-bound passkeys.
PKI-based Authenticators	Security Keys	Separate physical devices that store passkeys and connect to computers or mobile devices through USB, NFC, or BLE.
	Smart Cards	Physical cards containing an embedded security chip that stores PKI-based credentials and must be inserted into or connected to a device.
	Software Authenticators	Software-based authenticators that store PKI certificates.

Note: While this table covers the most common phishing-resistant MFA solutions used in enterprise environments, it is not exhaustive.

The Center thanks the following individuals for contributing their expertise:

Patrick Thielen - Liberty Mutual

Jud Dressler - Resilience

Matt McCabe & Scott Stransky - Marsh

Roopesh Shah - MunichRe

Craig Guiliano - Chubb

Mark Camillo - CyberAcuView

Matthieu Chan Tsin - Cowbell